

VEG: AN INTELLIGENT WORKBENCH FOR ANALYSING SPECTRAL REFLECTANCE DATA

P. Ann Harrison
JJM Systems, Inc.
1225 Jefferson Davis Hwy., Suite 412
Arlington, VA 22202
Tel: (703) 416-8256
FAX: (703) 416-8259

Patrick R. Harrison
U. S. Naval Academy

Daniel S. Kimes
NASA/GSFC Code 923

Abstract

An Intelligent Workbench (VEG) has been developed for the systematic study of remotely sensed optical data from vegetation. A goal of the remote sensing community is to infer physical and biological properties of vegetation cover (e.g. cover type, hemispherical reflectance, ground cover, leaf area index, biomass and photosynthetic capacity) using directional spectral data. Numerous techniques that infer some of these vegetation properties have been published in the literature. A fundamental problem is deciding which technique to apply to the data and then estimating the error bounds on the results. Studies have found that using conventional techniques produced errors as high as 45%.

VEG collects together in a common format techniques previously available from many different sources in a variety of formats. The decision as to when a particular technique should be applied is non-algorithmic and requires expert knowledge. VEG has codified this expert knowledge into a rule-based decision component for determining which technique to use. VEG provides a comprehensive interface that makes applying the techniques simple and aids a researcher in developing and testing new techniques. VEG also allows the scientist to incorporate historical databases into problem solving. The scientist can match the target data being studied with historical data so the historical data can be used to provide the coefficients needed for applying analysis techniques. The historical data also provides the basis for much more

accurate error estimates than were previously available. VEG also enables the scientist to try "what-if" experiments on data using a variety of different techniques and historical data sets to do comparative studies or test experimental hypotheses.

VEG also provides a classification algorithm that can learn new classes of surface features. The learning system uses the database of historical cover types to learn class descriptions of one or more classes of cover types. These classes can include broad classes such as soil or vegetation or more specific classes such as forest, grass and wheat. The classes can also include subclasses based on continuous parameters, e.g. 0-30% ground cover. The learning system uses sets of positive and negative examples from the historical database to find the most important features that uniquely distinguish each class. The system then uses the learned classes to classify an unknown sample by finding the class that best matches the unknown cover type data. The learning system also includes an option that allows the user to test the system's classification performance.

VEG was developed using object oriented programming, and the current version consists of over 1500 objects.

Introduction

The intent of this paper is to describe the advanced and novel concepts and features of the VEG system, and to show how VEG contributes to and extends the capabilities of

the scientist. VEG is an intelligent workbench for doing scientific studies of the earth's vegetation using optical reflectance data from sensor platforms. The system is being developed as a NASA/GSFC effort in the Biospherical Sciences branch. The workbench represents the development of a concept originally proposed on a much smaller scale by Abelson and Sussman (1987). Their workbench was intended to provide a tool that integrated a diverse set of concepts into an expressive environment for conducting scientific investigations. The VEG system provides a new and sophisticated intelligent system for the support of analysing spectral reflectance data of vegetation.

Background

The remote sensing community studies spectral data from the Earth's surface to infer physical and biological properties of vegetation. Large quantities of sensor data are collected and integrated to produce knowledge about surface characteristics such as cover type, ground cover, leaf area index, biomass and photosynthetic capacity. Future work using the Earth Observing System (EOS Reference Handbook, 1993) will produce significantly more complex as well as larger volumes of data. Studies of spectral reflectance data contribute critically important ecological information to a variety of scientific work including the effect of forest and natural vegetation clearing on local and regional climates, the relation of vegetation properties to energy and water balance, the relation between environmental parameters governing the energy balance and drought and desertification, and the relation between the absorbed, photosynthetically active radiation and the potential productivity of vegetation systems. The importance of these studies is discussed in detail in Kimes, Sellers and Newcomb (1987).

A central process in analysis is the application of a variety of extraction techniques to the raw spectral data to extract additional information for inferring surface characteristics. The fundamental problem is deciding which techniques to apply to the

data, and estimating the error bounds on the results. Studies have found that using traditional, ad hoc approaches, the errors of estimation were as high as 45% (proportion of true value) (Kimes, Harrison & Ratcliffe, 1991; Kimes and Sellers, 1985). Heuristic approaches, promise to overcome the simplicity and lack of flexibility of traditional algorithmic approaches and reduce estimation error by taking advantage of partial knowledge to make decisions about technique choice.

The basic datum being analyzed is directional optical reflectance data. Directional reflectance observations are made and then extraction techniques are used to relate these measurements to vegetation characteristics. Reflectance data can be collected on the ground, from aircraft or from satellites. The nature of this data is such that many decisions as to how to handle a particular data set need to be made at the expert level. The process of analysis is also complex and time consuming, requiring numerous steps and the comparison of new data with a potentially very large database of historical data with known attributes. The VEG workbench was designed to manage these problems.

Overview of VEG

VEG collects in a common format various techniques previously available in a hodgepodge of formats from a variety of different sources. VEG makes these techniques readily available to the scientist in one program. It also provides a rule-based decision tool for determining which technique to choose. It captures expertise in rules about when to use each technique. It captures the priority that should be given to different techniques by a simple weighting scheme. VEG provides a comprehensive interface that makes applying the techniques simple. VEG also incorporates historical databases into the problem solving process, enabling the matching of a target being studied to similar historical data so the historical data can be used to provide the coefficients needed for applying the techniques. The historical data also provides a much more accurate error

estimate than was previously available. VEG provides an interface for entering data from external files and outputting results to files in a variety of different formats. VEG also includes a toolbox which allows the user to browse the system, dynamically plot data, get help and print screen dumps.

The current version of VEG implements three different capabilities: estimation of vegetation parameters, estimation of atmospheric effects and a classification learning system. These capabilities represent the three subgoal categories in the system. The subgoal category "vegetation parameter techniques" enables the scientist to apply various techniques to calculate the surface properties, spectral hemispherical reflectance, total hemispherical reflectance, view angle extension and proportion ground cover. Subgoals in the category "atmospheric techniques" make atmospheric corrections to data. "Atmospheric techniques" allow satellites and aircraft data to be corrected for atmospheric effect to determine what the

equivalent ground level measurements would have been. Additional atmospheric techniques allow data collected at ground level to be projected to different atmospheric heights. These atmospheric capabilities are currently being implemented. The "classification learning system" subgoals category enables VEG to learn class descriptions of different vegetation classes and then use the learned classes to classify an unknown sample. The "neural networks" subgoal category provides for analysis using neural or connectionist networks. It is not yet available. Figure 1 shows a decomposition of basic VEG system goals.

VEG was implemented using object oriented programming. The objects in the VEG knowledge base were arranged in a loosely defined hierarchy organized by the major components: databases, control methods, techniques, tools and rules. Within the components, objects are organized in abstraction hierarchies. Separate subclasses hold the objects required by the "estimate vegetation parameter" and "estimate

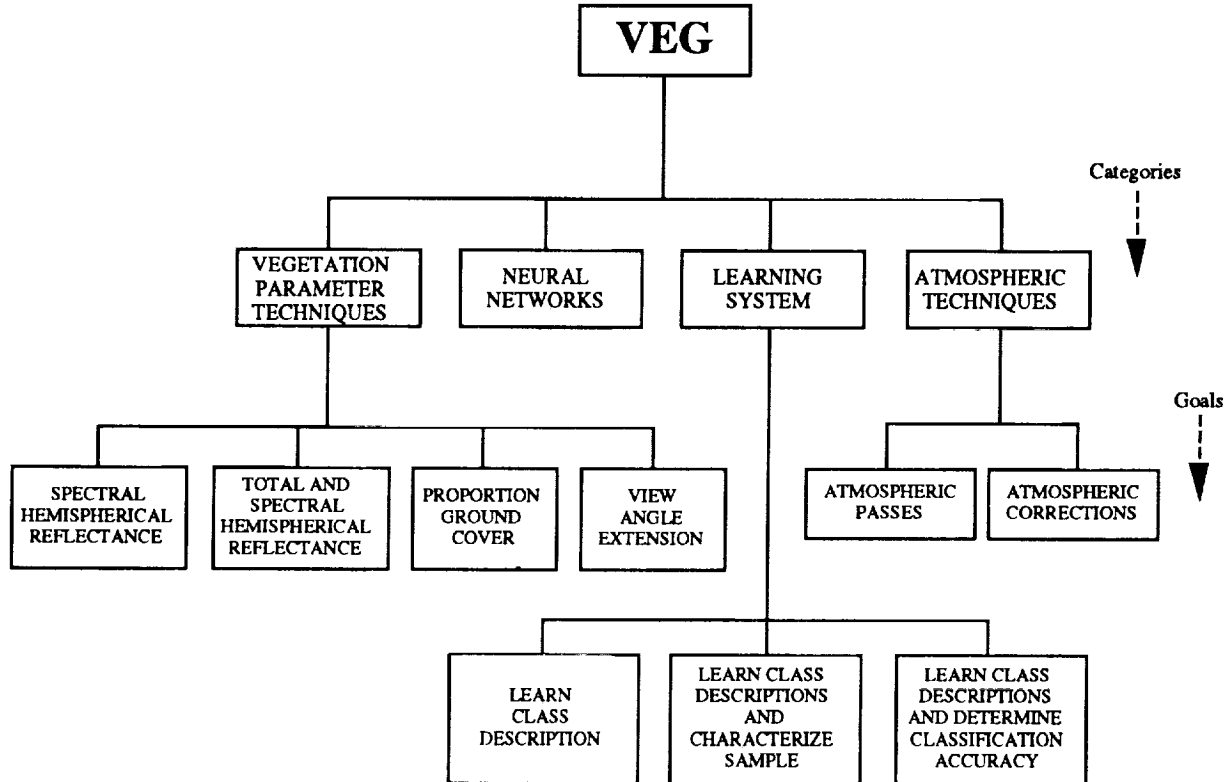


Figure 1: Goal Decomposition of VEG

atmospheric effect" goal categories. The learning system is housed in a separate knowledge base that is loaded only when needed. The full object system with data and rules loaded typically consists of about 1500 objects.

The database components of VEG include various databases used by the system. The most important database subclass contains various sets of typical cover type data which are used to test and demonstrate the VEG system. If VEG is run using new cover type data, additional units are constructed in this subclass to hold the new cover type data. During processing, additional objects are created to store the intermediate and final results of applying various techniques to a cover type sample. These can be inspected or browsed at any time.

All the options in VEG make use of the historical cover type database. This database contains results from experiments by scientists on a wide variety of different cover types. The historical cover type database is maintained externally. It is loaded when needed in a specific application. Currently this is in the form of cases stored in flat files. In the future, it is envisioned that a relational database environment will replace the flat files.

Some of the methods required by VEG are stored in objects. Other methods are stored in files external to the knowledgebase. When the VEG knowledgebase is loaded, these methods files are also loaded. The files contain compiled Common Lisp code for executing steps in processing data and applying the techniques.

Rules are used to determine which techniques to apply to a sample of cover type data. There is a different set of forward chaining rules for each VEG subgoal. In addition, the subgoal proportion ground cover has two sets of rules, one for single wavelength techniques, and one for multiple wavelength techniques. The rules are quite

complex. They combine execution of Common Lisp functions with traditional pattern matching. Figure 2 shows an example of a rule. This rule selects the technique 2FULL.1HALF.STRINGS if the data contains two full and one half strings.

VEG also contains a rulebase for ranking the techniques. Currently, the rules in this rulebase implement a simple weighting system. It is anticipated that a more complex rulebase for ranking techniques, incorporating more remote sensing expertise, will be added to VEG in the future.

The rules in VEG are all domain rules rather than control rules. System control is embedded in the window system through the ordering of windows and the constraints on the data input to any window.

VEG is embedded in an extensive, window-driven interface system that provides a variety of screens to enhance dialogue between the scientist and the system. The interface is a key feature of this system. It was designed to focus the scientist on the appropriate level of organization to carry out scientific work without attention to "housekeeping" functions. The interface allows the scientist to interact with VEG and select options at all stages of a run by clicking the mouse over the appropriate menu option. It prevents the user from selecting any step before the prerequisite steps have been carried out. The interface allows a scientist with no knowledge of Common Lisp or the detailed structure of VEG to use the system with ease.

Most operations are controlled using the mouse. The only time that the scientist needs to use the keyboard during a run is if he or she chooses to enter new data manually. When a new value is entered manually, a function is run. If the user has typed in an invalid value, a message is displayed and the value is not retained in the slot. Thus the interface provides validation of the input data. The interface also prevents incomplete data sets from being stored.

IF	(THE CURRENT.SAMPLE.WAVELENGTHS OF ESTIMATE.HEMISPHERICAL.REFLECTANCE IS ?X) (THE STRING.OBJECTS OF ?X IS ?NUM) (LISP (= (LENGTH ?NUM) 3)) (LISP (= 1 (COUNT-IF #'(LAMBDA (X) (EQ 'HALF (GET.VALUE X 'FULL))) ?NUM))) (LISP (= 2 (COUNT-IF #'(LAMBDA (X) (EQ 'FULL (GET.VALUE X 'FULL))) ?NUM)))
THEN	(LISP (ADD.VALUE ?X 'TECHNIQUES '2FULL.1HALF.STRINGS))))
In Plain Text:	
If	There is a unit containing data being studied at one wavelength. The unit contains data which can be characterized as containing 3 strings. Of these strings, one is a half string and two are full strings.
Then	Add the value 2FULL.1HALF.STRINGS to the TECHNIQUES slot of the unit.

Figure 2: The Rule that Selects the Technique 2FULL1HALF.STRINGS

An interface to an input file of unknown cover type data is available in VEG. The interface enables the user to name the input file and specify the format for the file. Using this format, the input file is read and the cover type data is stored for processing in the system. VEG also provides the user with the option of having the results of processing written to a file and selecting the format that should be used.

The toolbox is an important part of VEG. The user can activate the toolbox at any time during a run. The toolbox allows the user to read a description of the VEG system, browse the units and slots within the VEG system, obtain help about any screen, plot the zeniths, azimuths and reflectance values of reflectance data in two different plots, explore the historical data base and print out a screen dump of the current screen. The toolbox provides a means of managing the levels of abstraction the scientist sees and allows the scientist to deepen his understanding of system functionality.

A help system has also been integrated into VEG. The help system is

currently a prototype version of a system that would provide on-line help for a scientist using VEG. It would allow the scientist to get more information about each screen in the VEG interface. It was designed to help the new user of VEG to learn how to operate the system. Since the help system may not be needed by an experienced user, it was configured so that it is loaded only when needed. The first time the user asks for help, the help system is automatically loaded. An interface that allows the scientist to add and modify help messages has also been integrated into VEG. This enables the scientist to evolve the help system over time.

The Subgoal "Spectral Hemispherical Reflectance"

The steps in the subgoal "spectral hemispherical reflectance" are described in this section to illustrate how VEG can be used. When the option "spectral hemispherical reflectance" is selected, the menu shown in Figure 3 is displayed. This menu enables the user to invoke the steps involved in processing target data to estimate

the spectral hemispherical reflectance and estimate the error in the calculation. Before each step is carried out, a check is made to make sure that the prerequisite steps have been carried out. For example, the results cannot be output before the techniques have been executed. If any prerequisite steps have not been carried out, a message is displayed and the user is prompted to complete the prerequisite steps.

ENTER.DATA
CHARACTERIZE.INPUT
CHARACTERIZE.TARGET
CREATE.RESTRICTED.DATA
INTERP/EXTRAP.RESTRICTED.DATA
CHARACTERIZE.RESTRICTED.DATA
GENERATE.TECHNIQUES
RANK.TECHNIQUES
EXECUTE.TECHNIQUES
OUTPUT.RESULTS
SELECT.ALL.OPTIONS
INITIALIZE.SYSTEM
QUIT

**Figure 3: Steps in the Subgoal
"Spectral Hemispherical Reflectance"**

The first step is to enter the target data. The user can either enter a new, original set of data for an unknown target or select one of a number of samples of target data already stored in VEG. Each set of target data, whether entered by the user or selected from the samples already in VEG, can contain reflectance data at one or more wavelengths. Next, the target data at each wavelength is characterized. Sets of view angles in the same azimuthal plane are identified as "strings." Strings are characterized as full-strings if they contain both forwardscatter and backscatter data and half-strings if they contain either backscatter or forwardscatter data. Next the target is characterized. If the target data does not contain a value for ground cover or leaf area index, a crude estimation of these values is made in this step.

The next step is creating the restricted data set. This step involves selecting a subset of the historical database to be used for

generating the coefficients required by the techniques and estimating the error term when various techniques are applied to the target data. The selection of the restricted data set can either be made automatically by the system or manually by the user.

If the user elects to have the restricted data set selected automatically by the system, the database of historical cover types is searched to find the cover types that best match the target. The subset of historical cover types that matches the wavelength of the target is first identified. From this subset, the cover types whose ground cover and solar zenith angle are within ten percent of the values for the target are then identified and pushed onto a list. If the list contains insufficient values, the search is then widened to include cover types whose sun angles and proportion ground cover are within 20 percent of the values for target data. The search criteria are progressively widened until either sufficient cover types have been identified or all cover types whose sun angle and proportion ground cover are within 100 percent of the values in the target have been collected.

The user can also manually select the restricted data set. In this case, a screen is opened. This screen allows the user to enter the maximum and minimum values to be considered for parameters such as height and solar zenith angle. The database of historical cover types is searched to find the cover types that match the criteria entered by the user. The user can then select the matched cover types, enter new maximum and minimum values and match the data again or select a subset of the matched data.

Next, the raw reflectance data for each cover type in the restricted data set is interpolated and extrapolated so that the view angles exactly match at each wavelength the view angles in the target data. The data in the restricted historical data units are characterized using the same methods that were used to characterize the target.

Generating the techniques to be applied to the data is the next step. The techniques can be generated automatically or

selected by the user. If the user elects to have the system generate the techniques, rules are run and the techniques that are suitable for estimating the spectral hemispherical reflectance of the target are identified. If the user elects to choose the techniques manually, a screen containing the names of all the available spectral hemispherical reflectance techniques is opened. When the user left-clicks on the name of a technique, a brief description of the technique is displayed. A function is called to check whether the technique is suitable for the sample. If the technique is not suitable for the sample, an error message is displayed and the technique is deselected. Rules in the "rank techniques" rulebase are run next and the techniques are ranked according to a simple weighting scheme and then displayed in order. The user can select the best one, two or three techniques for each wavelength or pick all the selected techniques.

The techniques are applied to the data at each wavelength in the target. If a technique requires coefficients, the user is asked whether all or half the restricted data set should be used for generating the coefficients and estimating the error. The appropriate coefficient methods are applied as necessary. The techniques are applied to the restricted historical data and the difference between the calculated spectral hemispherical reflectance and the correct value for the spectral hemispherical reflectance stored in the database is calculated. Using the error measurements from several historical cover types, the root mean square error is calculated. This provides an estimate of the error involved in applying the technique to the target data.

In the final step, the results are displayed on the screen. For each technique, the estimate of the spectral hemispherical reflectance, the error estimates and the coefficients are displayed. The screen allows the user to flip between the results at different wavelengths. The user is then asked whether the results should be written to a file. The results for all the VEG subgoals, including the subgoal spectral hemispherical reflectance, can be written to a file.

The Learning System

The learning system provides a tool for classifying new data and for learning new classifications. The learning system uses historical data that represents positive and negative examples to learn classifications. The learned classifications can then be used to classify unknown samples. This is a form of supervised learning first discussed by Mitchell (1982). The theory upon which the learning system was based is discussed in detail in Kimes, Harrison and Harrison (1992).

The learning system provides the user with three different options. In Option 1, the system uses the database of historical cover type data to learn class descriptions of one or more classes of cover types. These classes can include broad classes such as soil or vegetation or more specific classes such as forest, grass or wheat. The classes can also include subclasses based on continuous parameters such as 0-30% ground cover, 31-70% ground cover and 71-100% ground cover. In Option 2, the system learns class descriptions for one or more classes and then uses the learned classes to classify an unknown sample by finding the class that best matches the unknown cover type data. Option 3 allows the user to test the system's classification performance. In this option, the system learns class descriptions for one or more classes and then classifies the appropriate samples in the data base. The percentage of correctly classified samples is then used to summarize the degree of classification accuracy achieved by the learning system.

The first step in Option 1 is to define the training problem. An interface allows the user to enter the solar zenith angle, wavelengths and directional view angles. In order to define the class whose description is to be learned, the user first selects a parameter. In the case of a continuous parameter such as ground cover, the range of possible values is displayed and the user is prompted to enter the maximum and minimum values for the class. In the case of a discrete parameter such as description, the screen displays the possible values of the

parameter and prompts the user to enter the value for the parameter in the class. For example, if the parameter is description, the class might be forest. VEG then checks the validity of the entered data and prompts the user to enter the data again if it is invalid. Additional class parameters can then be defined as necessary. For example, a class might be defined as forest with 70-100% ground cover. The user can then enter data for additional classes such as 31-70% ground cover.

The second step is for the system to learn the class descriptions for the classes that were defined in the previous step. The first step in learning the class descriptions is to generate the training sets. The system searches the historical cover type database and finds the cover types that best match the training problem. A cover type matches the training problem if it has data at all the wavelengths specified in the training problem, its solar zenith is close to the training problem solar zenith, and it has a value for every parameter specified in the class definition. Once a matching cover type has been identified, the values in the slots for each parameter in the class definition are examined. If the cover type data fits the class definition, the name of the cover type is added to the positive training set. Otherwise, it is added to the negative training set. In the first search through the data base, each matching cover type whose solar zenith is within 10% of the training problem's solar zenith is identified and added to the appropriate training set. If insufficient cover types have been found for the training sets, the search is then repeated. In the second search, matching cover types whose solar zenith is within 20% of the training problem solar zenith are identified. The process of increasing the bounds on the solar zenith and searching through the database is continued until either the positive or negative training set exceeds the maximum permissible size, both training sets exceed the minimum permissible size or the bounds have increased to $\pm 100\%$. The learning system is usually run with a minimum training set size of 8 units. If when the search ends either training set is found to be empty, a message is

displayed on the screen and the process of learning class descriptions is stopped.

Next, the raw reflectance data from the cover type data in the training sets at the appropriate wavelength is interpolated and extrapolated to match the view angles in the training problem at each wavelength.

Once the training sets have been set up, rules are run in order to determine the set of possible hypotheses that can be constructed for the data in each training set. The left-hand side of each rule tests the view angle data. If the rule fires, the appropriate Common Lisp function is called. Each function generates possible hypotheses to be used in the training problem.

For example, the rule LR.1 fires if the view angle data at a particular wavelength contains at least two view angles. The right-hand side of this rule calls the lisp function TRY-DIRECTION-RELATIONSHIPS which generates direction relationships for every possible pair of view angles in the data and adds these to the list of hypotheses to be tested on the training problem. An example of a direction relationship that might be generated by this function is,

(GREATER-THAN
0.64 (60 180) (30 180)).

This relationship represents the hypothesis that at wavelength 0.64 μm , the reflectance at the view angle (60 180) is greater than the reflectance at view angle (30 180).

When the forward chaining of the rules has been completed, the set of all possible separate hypotheses for each training problem has been generated.

The next step in learning the class descriptions is to determine the discrimination score for each separate hypothesis. Each hypothesis such as (GREATER-THAN 0.64 (60 180)(30 180)) is tested on each sample in the positive and negative training sets. The sample score is 1 if the hypothesis is true and 0 otherwise. The discrimination score is calculated as:

$$\left[\frac{1}{p} \sum_{i=1}^p S_i \right] - \left[\frac{1}{n} \sum_{j=1}^n S_j \right] \quad (1)$$

where each sample score is S , S_i is the i th positive sample score, S_j is the j th negative sample score, p is the number of samples in the positive training set and n is the number of samples in the negative training set. Thus a discrimination score of 1 for a hypothesis represents the case where the hypothesis is true for all samples in the positive training set and false for all samples in the negative training set. This represents perfect discrimination. A score of 0 is the break even point where there is no effective discrimination between the positive and negative training sets. A score of less than zero for a hypothesis represents the case where the hypothesis is true for more samples in the negative training set than in the positive training set. In this case, the converse of the hypothesis would yield a positive discrimination score. For each hypothesis such as (GREATER-THAN 0.64 (60 180)(30 180)) two separate scores are calculated. The order of the elements is re-ordered and two scores such as:

(((((GREATER-THAN
(60 180)(30 180)) T) 0.64)) 0.4) (2)

and

(((((GREATER-THAN
(60 180)(30 180)) NIL) 0.64)) -0.4) (3)

are reported. In this example, the score ((((((GREATER-THAN (60 180)(30 180)) T) 0.64)) 0.4) means that the hypothesis that the reflectance at angle (60 180) is greater than the reflectance at angle (30 180) for the wavelength 0.64 μm produced a discrimination score of 0.4. The discrimination score in (2) is calculated directly by testing the hypothesis (GREATER-THAN 0.64 (60 180)(30 180)) on all the data in the positive and negative training sets. The discrimination score in (3), -0.4, is calculated as minus one multiplied by the discrimination score in (2). Scores such as (2) and (3) are calculated for each hypothesis.

The next step in the learning of class descriptions is to construct compound hypotheses. A compound hypothesis is composed of the combination of two or more individual hypotheses. The idea is that the interactions between various individual hypotheses may account for more variance (be more predictive) than any individual hypothesis. All the individual hypotheses are considered as potential parts of compound hypotheses, and not just the best single hypothesis.

Before compound hypotheses are constructed, heuristics are used to reduce the set of hypotheses for each training problem by removing any hypothesis that could not be combined with another hypothesis to form a compound hypothesis with a discrimination score better than the current best score. For this reason, every hypothesis whose positive training set score is less than or equal to the current best score for the problem is removed from the list of hypotheses. Hypotheses that do not discriminate or that score zero for the negative training set are also removed from the list of hypotheses. At the end of this step, the list of single hypotheses of each training problem contains only those hypotheses that could potentially be combined with other hypotheses to form a compound hypothesis with a discrimination score greater than the current best score for the problem.

The list of single hypotheses may contain in excess of fifty hypotheses, even after it has been reduced. The number of possible compound hypotheses for some training problems is immense. The problem of dealing with such a large number of potential compound hypotheses was the subject of much effort. Several alternative strategies were experimented with before a successful solution to the problem was found. The first attempt was to implement a breadth-first search. Compound hypotheses that had been investigated were stored on an explored list. Each time a compound hypothesis was investigated, all possible combinations of the hypothesis and other hypotheses were constructed and stored on an unexplored list. Checks were made to prevent duplication of compound hypotheses

on the unexplored list and to prevent the same hypothesis from being investigated more than once. This involved sorting all the separate hypotheses within a compound hypothesis

best score. If the best discrimination score for a single hypothesis is equaled by a compound hypothesis, the compound hypothesis and score are not stored. Once the level 2 search has been completed, a

The flexibility of the system allows the scientist a platform to conduct any number of explorations of a large body of reflectance data in a very short period of time. What took days in the past can now be accomplished in minutes. This means that the scientist can be much more productive and expansive in his/her thinking than would have been allowable without the time contraction and complexity management that this system provides.

The learning system provides a tool for classifying new data and for learning new classifications. The learning system uses historical data that represents positive and negative examples to learn classifications. The learned classifications can then be used to classify unknown samples. This is a form of supervised learning.

VEG was developed as a NASA/GSFC effort in the Biospherical Sciences branch. It is now being used by remote sensing scientists. It has proved to be a highly useful tool supporting scientific investigation as described by Kimes, Harrison and Ratcliffe (1991), Kimes and Holben (1992), Kimes, Harrison and Harrison (1992), Kimes, Irons and Levine (1992), Kimes and Deering (1992), Kimes, Kerber and Sellers (1993), and Kimes, Harrison and Harrison (1994).

Acknowledgment

This work was funded by the NASA Remote Sensing Science Program in Code SEP03 at NASA Headquarters.

References

Abelson, H. and Sussman, G.J. (May 1987). The Dynamicist's Workbench: 1 Automatic Preparation of Numerical Experiments. *MIT AI Memo 955*.

Kimes, D. S. and Deering, D. W. (1992). Remote Sensing of Surface Hemispherical Reflectance (Albedo) Using Pointable Multispectral Imaging Spectrometers. *Remote Sensing of Environment*, 39, 85-94.

Kimes, D.S., Harrison, P.R. and Ratcliffe, P.A. (October 1991). A knowledge-based expert system for inferring vegetation characteristics. *International Journal Remote Sensing*, 12(10), 1987-2020.

Kimes, D. S., Harrison, P. R., and Harrison P. A. (1994). Extension of Off-Nadir View Angles for Directional Sensor Systems. *Remote Sensing of Environment*, (Submitted).

Kimes, D.S., Harrison, P.R. and Harrison, P.A. (March 1992). Learning Class Descriptions from a Data Base of Spectral Reflectance with Multiple View Angles. *IEEE Transactions on Geoscience and Remote Sensing*, 30(2), 315-325.

Kimes, D. S. and Holben, B. N. (1992). Extracting Spectral Albedo from NOAA 9 AVHRR Multiple View Data Using an Atmospheric Correction Procedure and an Expert System. *Remote Sensing of Environment*, 13, 275-289.

Kimes, D. S., Irons, J. R., and Levine, E. R. (1993). Learning Class Descriptions from a Data Base of Spectral Reflectance of Soil Samples. *Remote Sensing of Environment*, 43, 161-169.

Kimes, D. S., Kerber, A. G., and Sellers, P. A. (1993). Spatial Errors in Creating Hemispherical Reflectance (Albedo) Maps from Directional Reflectance Data. *Remote Sensing of Environment*, 45, 85-94.

Kimes, D.S. and Sellers, P.J. (1985). Inferring hemispherical reflectance of the Earth's surface for global energy budgets from remotely sensed nadir or directional radiance values. *Remote Sensing of*

Reflectance Variations of Vegetation Canopies and Implications for Global and Regional Energy Budget Studies. *Journal of Climate and Applied Meteorology*, 26(8), 957-972.

Mitchell, T.M. (1982). Generalization as Search. *AI*, 18, 203-226.

Wharton, S.W. (1987). A Spectral-Knowledge-based Approach for Urban Landcover Discrimination. *IEEE Transactions on Geoscience and Remote Sensing*, GE-25, 272-282.

